

Publicly Verifiable Outsourced Computation with a Key Distribution Centre

James Alderman*, Carlos Cid**, Jason Crampton***, and Christian Janson†

Information Security Group, Royal Holloway, University of London

Abstract. The combination of software-as-a-service and the increasing use of mobile devices gives rise to a considerable difference in computational power between servers and clients. Thus, there is a desire for clients to outsource the evaluation of complex functions to a server and to be able to verify that the resulting value is correct. Previous work in this area of Publicly Verifiable Outsourced Computation (PVC) requires a costly pre-processing stage. However, in many practical situations multiple clients will be interested in the same set of core functions and will make use of the same servers. Thus, the pre-processing phase may be performed many more times than is necessary. In this paper we introduce a Key Distribution Center (KDC) that handles the generation and distribution of the keys that are required to support PVC, thereby eliminating this redundancy. We define a number of new security models and functionalities that arise with the introduction of the KDC, and present a construction of such a scheme built upon Key-Policy Attribute-based Encryption.

Keywords— Publicly Verifiable Outsourced Computation, Key Distribution Center, Key-policy Attribute-based Encryption, Revocation

1 Introduction

It is increasingly common for mobile devices to be used as general computing devices. There is also an increasing trend towards cloud computing and enormous volumes of data (“big data”) which mean that computations may require considerable computing resources. In short, there is, increasingly, a discrepancy between the computing resources of end-user devices and the resources required to perform complex computations on large datasets. This discrepancy, coupled with the increasing use of software-as-a-service, means there is a requirement for a client device to be able to delegate a computation to a server.

* James.Alderman.2011@live.rhul.ac.uk

** Carlos.Cid@rhul.ac.uk

*** Jason.Crampton@rhul.ac.uk

† Christian.Janson.2012@live.rhul.ac.uk

Consider, for example, a company that operates a “bring your own device” policy, enabling employees to use smartphones and tablets. It may not be possible for these devices to perform complex computations locally. Instead, a computation is outsourced over some network to a more powerful server (possibly outside the company, offering software-as-a-service, and hence untrusted) and the result of the computation is returned to the client device. Another example scenario arises in the context of battlefield communications where each member of a squadron of soldiers is deployed with a reasonably light-weight computing device. The soldiers gather data from their surroundings and send it to regional servers for analysis before receiving tactical commands based on results. Thus a soldier must have an assurance that the command has been computed correctly and by a trusted party. A final example could consider sensor networks where lightweight sensors transmit readings to a more powerful base station to compute statistics that can be verified by an experimenter.

In simple terms, given a function F to be computed by a server S , the client sends input x to S , who should return $F(x)$ to the client. However, there may be an incentive for the server (or an imposter) to cheat and return an invalid result $y \neq F(x)$ to the client. The acceptance of an incorrect result may have an advantage for the server, or the server may be too busy or may not wish to devote resources to perform the computation. Thus, the client wishes to have some assurance that the result y returned by the server is, in fact, $F(x)$. This problem, known as *Verifiable Outsourced Computation* (VC), has attracted a lot of attention in the community recently (see Sect. 2 for a brief overview). Many current schemes have an expensive pre-processing stage run by the client, which should be amortised over many function evaluations over distinct inputs. Note, however, it is likely that many different clients will be interested in outsourcing computations, and also that the functions of interest to each of the clients will substantially overlap, as in the “bring your own device” scenario discussed above. It is also conceivable that the number of computation servers offering to perform such computations will be relatively low (limited to a reasonably small number of trusted companies with plentiful resources). Thus, it is easy to envisage a situation in which many computationally limited clients wish to outsource the computation of the same function F to the same server, yet each must individually expend considerable resources to run the setup phase.

Our main contribution in this paper is to introduce a Key Distribution Center (KDC), that is responsible for running the setup stage on behalf of *all* clients. Thus, the expensive algorithm is executed just once and by the more capable KDC, rather than multiple times by restricted client devices. We consider two example settings: one is a straightforward generalisation of the previously considered model where clients send computations directly to an available server; in the second setting, we allow a pool of computational servers governed by some managing entity. Clients submit jobs to this pool and the manager distributes work according to a scheduling policy or a bidding process, and the result is returned to the client – thus the client may not require knowledge of the server identity or credentials beforehand.

We give definitions for a new framework of Publicly Verifiable Outsourced Computation that both removes redundancy and facilitates additional functionality (such as revoking misbehaving servers), including several new security notions. We also give a provably secure instantiation that meets the new definitions. In the manager model, we allow for “blind verification” by the manager or other entities, a form of output privacy, such that he learns whether the result is valid but not the value of the output. Thus he may reward or punish servers appropriately without learning function outputs.

It may be tempting to suggest that the KDC, as a trusted entity, performs all computations itself. However we believe that this is not a practical solution in many real world scenarios, e.g. the KDC could be an authority within the organisation responsible for user authorisation that wishes to enable workers to securely utilise cloud-based software-as-a-service. As an entity within the boundaries of the organisation, performing all computations would negate the benefits gained from outsourcing computations to externally available powerful servers. Additionally, as an authority on users and keys, the KDC may have simultaneous responsibilities in other systems, and we minimise its workload to key generation and revocation only.

2 Verifiable Computation Schemes and Related Work

The concept of non-interactive verifiable computation was introduced by Genaro et al. [4] and may be seen as a protocol between two polynomial-time parties, a *client*, C , and a *server*, S . A successful run of the protocol results in the provably correct computation of $F(x)$ by the server for an input x supplied by the client. More specifically, a VC scheme comprises the following steps [4]:

1. **KeyGen** (*Run once*): C computes evaluation information EK_F that is given to S to enable it to compute F
2. **ProbGen** (*Run multiple times*): C sends the encoded input σ_x to S
3. **Compute** (*Run multiple times*): S computes $y = F(x)$ using EK_F and σ_x and returns an encoding of the output σ_y to C
4. **Verify** (*Run multiple times*): C checks whether σ_y encodes $F(x)$

Parno et al. [8] introduced the idea of *Publicly Verifiable Computation* (PVC). In this setting, a single client C_1 computes EK_F , as well as publishing information PK_F that enables other clients to encode inputs, meaning that only one client has to run the expensive pre-processing stage. Each time a client submits an input x to the server, the client may publish $VK_{F,x}$, which enables any other client to verify that the output is correct. A PVC scheme uses the same four algorithms as VC but **KeyGen** and **ProbGen** are now required to output public values that other clients may use to encode inputs and verify outputs, respectively.

PVC using KP-ABE. Parno et al. provide a concrete instantiation using *Key-policy Attribute-based Encryption*¹ (KP-ABE) [7], for Boolean functions [8].

¹ If input privacy is required then a predicate encryption scheme could be used in place of the KP-ABE scheme.

Table 1: PVC using KP-ABE

Abstract PVC parameter	Parameter in KP-ABE instantiation
EK_F	$SK_{\mathbb{A}_F}$
PK_F	Master public key PP
σ_x	Encryption of m using PP and A_x
σ_y	m or \perp
$VK_{F,x}$	$g(m)$

Define a universe \mathcal{U} of n attributes and associate $V \subseteq \mathcal{U}$ with a binary n -tuple in which the i th place is 1 if and only if the i th attribute is in V . We call this the *characteristic tuple* of V . Thus, there is a natural one-to-one correspondence between n -tuples and attribute sets; we write A_x to denote the set associated with x . A function $F : \{0, 1\}^n \rightarrow \{0, 1\}$ is monotonic if $x \leq y$ implies $F(x) \leq F(y)$, where $x = (x_1, \dots, x_n)$ is less than or equal to $y = (y_1, \dots, y_n)$ if and only if $x_i \leq y_i$ for all i . For a monotonic function F , the set $\{x \in \{0, 1\}^n : F(x) = 1\}$ defines a monotonic access structure which we denote \mathbb{A}_F .

The mapping between PVC and KP-ABE parameters are shown in Table 1. Informally, for a Boolean function F , the client generates a private key $SK_{\mathbb{A}_F}$ using the KP-ABE KeyGen algorithm. Given an input x , a client encrypts a random message m “with” A_x using the KP-ABE Encrypt algorithm and publishes $VK_{F,x} = g(m)$ where g is a suitable one-way function (e.g. a hash function). The server decrypts the message using the KP-ABE Decrypt algorithm, which will either return m (when $F(x) = 1$) or \perp . The server returns m to the client. Any client can test whether the value returned by the server is equal to $g(m)$. Note, however, that a “rational” malicious server will always return \perp , since returning any other value will (with very high probability) result in the verification algorithm returning a reject decision. Thus, it is necessary to have the server compute both F and its “complement” (and for both outputs to be verified). We revisit this point in Appendix A. The interested reader may also consult the original paper for further details [8]. Note that in order to compute the private key $SK_{\mathbb{A}_F}$, it is necessary to identify all minimal elements x of $\{0, 1\}^n$ such that $F(x) = 1$. There may be exponentially many such x . Thus, the initial phase is indeed computationally expensive for the client. Note also that the client may generate different private keys to enable the evaluation of different functions.

Other Related Work. The concept of *non-interactive* verifiable computation was formalised by Gennaro et al. [4] who gave a construction using Garbled Circuits [9] with fully homomorphic encryption [5] to re-randomise the circuit to allow multiple executions. In independent and concurrent work, Carter et al. [2] introduce a third party to generate garbled circuits for such schemes but require this entity to be online throughout and model the system as a secure multi-party computation between the client, server and third-party. Some works [6, 3] consider the multi-client case where functions are computed over joint input from multiple clients and notions such as input privacy become more important.

Algorithm	Run by		
	VC	PVC	PVC-KDC
KeyGen	C_1	C_1	KDC
ProbGen	C_1	C_1, C_2, \dots	C_1, C_2, \dots
Compute	S	S	S_1, S_2, \dots
Verify	C_1	C_1, C_2, \dots	C_1, C_2, \dots or M

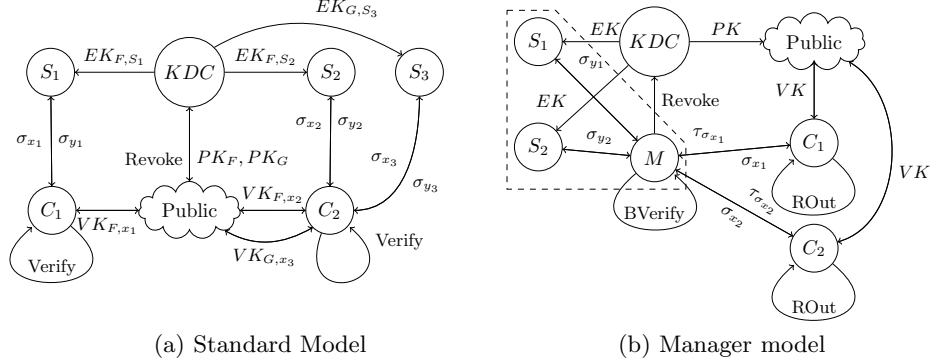


Fig. 1: The operation of PVC-KDC

3 PVC with a Key Distribution Center

We now introduce our extension of PVC, which we call *Publicly Verifiable Computation with a Key Distribution Center* (PVC-KDC). We assume there are many clients and multiple servers. Different servers may compute the same function F and servers are “certified” to compute F by the Key Distribution Center. As we briefly explained in the introduction, there appear to be good reasons for adopting an architecture of this nature and several scenarios in which such an architecture would be appropriate. The increasing popularity of relatively lightweight mobile computing devices in the workplace means that complex computations may best be performed by more powerful servers run by the organization or in the cloud and we would wish to have some guarantee that those servers are certified to perform certain functions. It is essential that we can verify the results of the computation. If cloud services are competing on price to provide “computation-as-a-service” then it is important that a server cannot obtain an unfair advantage by simply not bothering to compute $F(x)$ and returning garbage instead. It is also important that a server who is not certified cannot return a result without being detected. In this paper we focus on two example system architectures, which we call the Standard Model and the Manager Model:

- The standard model is a natural extension of the PVC architecture with the addition of a KDC. The entities comprise a set of clients, a set of servers and a KDC. The KDC initialises the system and generates keys to enable verifiable computation. Clients submit computation requests to a particular server and publish some verification information. Any party can verify the correctness of a server’s output. If the output is incorrect, the client may

report the server to the the KDC for revocation which will prevent the server from performing any further computations of this function.

- The manager model, in contrast, employs an additional Manager entity who “owns” a pool of computation servers. Clients submit jobs to the manager, who will select a server from the pool based on workload scheduling, available resources or as a result of some bidding process if servers are to be rewarded per computation. A plausible scenario is that servers enlist with a manager to “sell” the use of spare resources, whilst clients subscribe to utilise these through the manager. Results are returned to the manager who should be able to verify the server’s work. The manager forwards correct results to the client whilst a misbehaving server may be reported to the KDC for revocation, and the job assigned to another server. Due to public verifiability, any party with access to the output and the verification token can also verify the result. However, in many situations we may not desire external entities to access the result, yet there remains legitimate reasons for the manager to perform verification. Thus we introduce “blind verification” such that the manager (or other entity) may verify the validity of the computation without learning the output, but the delegating client holds an extra piece of information that enables the output to be retrieved.

A PVC-KDC system operates as follows, and as shown in Figure 1:

1. The KDC initializes the system and generates public and private parameters, which are used to generate new keys for many different functions.
2. A server S may join the system by registering with the KDC to receive a private key SK_S .
3. A server S that wishes to provide a computation service for a function F makes a request to the KDC, who generates a (personalised) *secret* value $EK_{F,S}$ and transmits it to S ; this will be used in the computation of $F(x)$.²
4. The KDC also generates public data PK_F for F , used to encode client inputs, and publishes a list of servers, L_F , that are certified to compute F .
5. To outsource the computation of $F(x)$, a client C uses PK_F to prepare σ_x and public verification value $VK_{F,x}$. In the standard model, the client sends σ_x to a selected server; in the manager model, σ_x is sent to a manager who distributes it to a server according to some policy or bidding process.
6. On receipt of a request to evaluate $F(x)$, (an honest server) S computes the encoded output σ_y using $EK_{F,S}$ and σ_x , publishing or returning σ_y to C in the standard model, or to the manager in the manager model.
7. In the standard model, any party, V , can run a verification algorithm using σ_y and $VK_{F,x}$ as inputs; the algorithm will output **accept** and $F(x)$ if and only if σ_y encodes $F(x)$. In the manager model, the manager (or other) runs a blind verification algorithm such that they learn whether the output is valid, but *not* the actual value of $F(x)$ – thus the manager learns whether to reward or request the revocation of the server and may enlist an additional server if required, but does not learn the (sensitive) result. The delegator

² Note that for the purposes of revocation $EK_{F,S}$ is associated to a particular server S and is not public as in previous schemes.

or other chosen verifiers hold additional information that can be used in an output retrieval algorithm to learn the value of $F(x)$ without having to verify the result and potentially resubmit. These steps could be run together as in the standard model, and we collectively call these the verification algorithm.

8. If an invalid result is detected, the verifier or manager reports S by sending a token τ_{σ_y} to the KDC, who will revoke S . Thus S may incur a financial penalty from being unable to compute F until the KDC re-certifies him.

3.1 Formal Details

Definition 1. A *Publicly Verifiable Outsourced Computation Scheme with Key Distribution Center (PVC-KDC)* comprises the following algorithms:

- **Setup**(1^λ) \rightarrow (PP, MK): Run by the KDC to establish public parameters PP and a master secret key MK.
- **FnInit**(PP, MK, F) \rightarrow (PK_F, L_F): Run by the KDC to generate a public delegation key, PK_F , for a function F as well as a list L_F of available servers for evaluating F , which is initially empty.
- **Register**(PP, MK, S) $\rightarrow SK_S$: Run by the KDC to generate a personalised key SK_S for a computation server S .
- **Certify**(PP, MK, F, L_F, S) $\rightarrow (EK_{F,S}, L_F)$: Run by the KDC to generate a certificate in the form of an evaluation key $EK_{F,S}$ for a function F and server S . S is added to the list, L_F , of available servers for evaluating F .
- **ProbGen**(x, PK_F) $\rightarrow (\sigma_x, VK_{F,x}, b)$: The **ProbGen** algorithm is run by a client to delegate the computation of $F(x)$ to a server. The output value b is used to enable output retrieval after the blind verification step.
- **Compute**($\sigma_x, EK_{F,S}, SK_S$) $\rightarrow \sigma_y$: Run by a server S in possession of an evaluation key $EK_{F,S}$, SK_S and an encoded input σ_x of x to evaluate $F(x)$ and output an encoding, σ_y , of the result, which includes an identifier of S .
- **Verify**(PP, $\sigma_y, VK_{F,x}, L_F$) $\rightarrow (\tilde{y}, \tau_{\sigma_y})$: Verification consists of two steps.
 - **BlindVerify**(PP, $\sigma_y, VK_{F,x}, L_F$) $\rightarrow (\mu, \tau_{\sigma_y})$: Run by any verifying party (standard model), or run by the manager (manager model), in possession of $VK_{F,x}$ and encoded output, σ_y . This outputs a token $\tau_{\sigma_y} = (\text{accept}, S)$ if the output is valid, or $\tau_{\sigma_y} = (\text{reject}, S)$ if S misbehaved. It also outputs μ which is an encoding of the actual output value.
 - **RetrieveOutput**($\mu, \tau_{\sigma_y}, VK_{F,x}, b$) $\rightarrow \tilde{y}$: Run by a verifier in possession of b to retrieve the actual result \tilde{y} which is either $F(x)$ or \perp .
- **Revoke**(MK, τ_{σ_y}, F, L_F) $\rightarrow (\{EK_{F,S'}\}, L_F)$ or \perp : Run by the KDC if a misbehaving server is reported i.e. that **Verify** returned $\tau_{\sigma_y} = (\text{reject}, S)$ (if $\tau_{\sigma_y} = (\text{accept}, S)$ then this algorithm should output \perp). It revokes the evaluation key $EK_{F,S}$ of the server S thereby preventing any further evaluations of F . This is achieved by removing S from L_F (the list of servers for F) and issuing updated evaluation keys $EK_{F,S'}$ to all servers $S' \neq S$.

Definition 2 (Correctness). A *Publicly Verifiable Computation Scheme with a Key Distribution Center (PVC-KDC)* is correct for a family of functions \mathcal{F} if

for all functions $F \in \mathcal{F}$ and inputs x , where $\text{negl}(\cdot)$ is a negligible function of its input:

$$\begin{aligned} & \Pr[\text{Setup}(1^\lambda) \rightarrow (\text{PP}, \text{MK}), \text{FnlNit}(\text{PP}, \text{MK}, F) \rightarrow (PK_F, L_F), \\ & \quad \text{Register}(\text{PP}, \text{MK}, S) \rightarrow SK_S, \text{Certify}(\text{PP}, \text{MK}, F, L_F, S) \rightarrow (EK_{F,S}, L_F), \\ & \quad \text{ProbGen}(x, PK_F) \rightarrow (\sigma_x, VK_{F,x}, b), \\ & \quad \text{Verify}(\text{PP}, \text{Compute}(\sigma_x, EK_{F,S}, SK_S), VK_{F,x}, L_F) \rightarrow (F(x), (\text{accept}, S))] \\ & = 1 - \text{negl}(\lambda). \end{aligned}$$

3.2 Security Models

We now formalise several notions of security as a series of cryptographic games. The adversary against a particular function F is modelled as a probabilistic polynomial time algorithm \mathcal{A} run by a challenger. The adversary algorithm may maintain state and be multi-stage and we overload the notation by calling each of these adversary algorithms \mathcal{A} . The notation \mathcal{A}^O is used to denote the adversary \mathcal{A} being provided with oracle access to the following functions: $\text{FnlNit}(\text{PP}, \text{MK}, \cdot)$, $\text{Register}(\text{PP}, \text{MK}, \cdot)$, $\text{Certify}(\text{PP}, \text{MK}, \cdot, \cdot, \cdot)$ and $\text{Revoke}(\text{MK}, \cdot, \cdot, \cdot)$. In each of the games, we define the *advantage* and *security* of \mathcal{A} as:

Definition 3. The advantage of an adversary \mathcal{A} running in probabilistic polynomial time (PPT), making a polynomial number of queries q is defined as follows, where $\mathbf{X} \in \{\text{PubVerif}, \text{Revocation}, \text{VindictiveS}, \text{BVerif}, \text{VindictiveM}\}$:

$$\text{Adv}_{\mathcal{A}}^{\mathbf{X}}(\mathcal{PVC}_{KDC}, F, 1^\lambda, q) = \Pr[\mathbf{Exp}_{\mathcal{A}}^{\mathbf{X}}[\mathcal{PVC}_{KDC}, F, 1^\lambda] = 1].$$

A PVC-KDC is secure against Game \mathbf{X} for a function F , if for all PPT adversaries \mathcal{A} , $\text{Adv}_{\mathcal{A}}^{\mathbf{X}}(\mathcal{PVC}_{KDC}, F, 1^\lambda, q) \leq \text{negl}(\lambda)$.

Public Verifiability. In Game 1 we extend the Public Verifiability game of Parno et al. [8] to formalize that multiple servers should not be able to collude to convince *any* verifying party of an incorrect output (i.e. that **Verify** returns **accept** on a σ_y for $y \neq F(x)$). The game begins (line 1) with the adversary selecting a (polynomially sized) set of n input values that he would like to see the problem encoding of (and the corresponding time period³). The challenger runs **Setup**, **FnlNit** and **Register** to initialise the system and create a public delegation key for a function F given as a parameter to the game (lines 2 to 4). The adversary is given the delegation key, his private key and the public parameters (i.e. all values known to a server in the real setting), and must output a list of servers that should be certified to compute F (line 6)⁴.

³ The time period here is changed every time a server is revoked. Alternatively, the time period could be regularly updated but the **Revoke** algorithm must be run at each interval even if the revocation list has not changed.

⁴ This corresponds to the revocation list in the model of [1] except that we consider a certification list of servers that should receive the update keys rather than a revo-

Game 1 $\text{Exp}_{\mathcal{A}}^{\text{PubVerif}}[\mathcal{PVC}_{\text{KDC}}, F, 1^\lambda]:$

```

1:  $\{t_i^*, x_i^*\}_{i \in [n]} \leftarrow \mathcal{A}(1^\lambda);$ 
2:  $(\text{PP}, \text{MK}) \leftarrow \text{Setup}(1^\lambda);$ 
3:  $(PK_F, L_F) \leftarrow \text{FnInit}(\text{PP}, \text{MK}, F);$ 
4:  $SK_{\mathcal{A}} \leftarrow \text{Register}(\text{PP}, \text{MK}, \mathcal{A});$ 
5:  $L_F \leftarrow \mathcal{A}(PK_F, \text{PP}, L_F, SK_{\mathcal{A}});$ 
6:  $(EK_{F, \mathcal{A}}, L_F) \leftarrow \text{Certify}(\text{PP}, \text{MK}, F, L_F, \mathcal{A});$ 
7: for  $i = 1$  to  $n$  do
8:    $\{\sigma_{x_i^*}, VK_{F, x_i^*}, b_i\} \leftarrow \text{ProbGen}(\{t_i^*, x_i^*\}, PK_F);$ 
9:    $\sigma_{y^*} \leftarrow \mathcal{A}^\mathcal{O}(PK_F, \text{PP}, L_F, \{\sigma_{x_i^*}, VK_{F, x_i^*}\}, EK_{F, \mathcal{A}}, SK_{\mathcal{A}});$ 
10: if  $\exists i \in [n]$  s.t.  $((\tilde{y}, \tau_{\sigma_{y^*}}) \leftarrow \text{Verify}(\text{PP}, \sigma_{y^*}, VK_{F, x_i^*}, L_F))$  and  $((\tilde{y}, \tau_{\sigma_{y^*}}) \neq (\perp, (\text{reject}, \mathcal{A})))$  and  $(\tilde{y} \neq F(x_i^*))$  then
11:   return 1
12: else
13:   return 0

```

The challenger then runs **ProbGen** for each challenge input and gives the encoded inputs to the adversary. The adversary also has oracle access to model the corruption of other servers (line 10), and aims to create an encoded output that is accepted by the challenger yet is not valid for any challenge input.

Revocation. In Game 2 we require that if a server is detected as misbehaving (i.e. **Verify** outputs $(\perp, (\text{reject}, S))$) then any subsequent evaluations of F by S should be rejected. Even though we have outsourced the costly computation and pre-processing stages to the server and KDC respectively, there is still a cost involved in delegating and verifying a computation. If a server is known not to be trustworthy then we remove any incentive for it to attempt to provide an outsourcing service (since it knows the result will not be accepted). In addition, we may like to punish malicious servers by removing their ability to perform work (and earn rewards) for a period of time. Finally, from a privacy perspective, we may not wish to supply input data to a server that is known not to be trustworthy. In this game the adversary chooses the target input values as before (line 1) but now the evaluation key $EK_{F, \mathcal{A}}$ that it had access to when selecting x is revoked (line 8) before the computation is run. We require that the adversary is no longer able to provide *any* result that verifies correctly (even $F(x)$).

Vindictive Servers. The motivation for this notion of security is the manager model where the client does not know the identities of servers selected from the pool. Now, since an invalid result can lead to revocation, this leads to a new threat model (particularly in systems where servers gain rewards per computation performed) in which a malicious server may return incorrect results but attribute them to an alternate server ID such that the (honest) server is revoked,

cation list of servers that should not receive these keys. The requirement to output this list here is due to the selective IND-sHRSS game that we base the construction upon. Since this is used in a black-box manner however, a stronger primitive may allow this game to be improved accordingly.

Game 2 $\text{Exp}_{\mathcal{A}}^{\text{Revocation}} [\mathcal{PVC}_{\text{KDC}}, F, 1^\lambda]$:

```

1:  $\{t_i^*, x_i^*\}_{i \in [n]} \leftarrow \mathcal{A}(1^\lambda)$ ;
2:  $(\text{PP}, \text{MK}) \leftarrow \text{Setup}(1^\lambda)$ ;
3:  $(PK_F, L_F) \leftarrow \text{FnInit}(\text{PP}, \text{MK}, F)$ ;
4:  $SK_{\mathcal{A}} \leftarrow \text{Register}(\text{PP}, \text{MK}, \mathcal{A})$ ;
5:  $L_F \leftarrow \mathcal{A}(PK_F, \text{PP}, L_F, SK_{\mathcal{A}})$ ;
6:  $(EK_{F,\mathcal{A}}, L_F) \leftarrow \text{Certify}(\text{PP}, \text{MK}, F, L_F, \mathcal{A})$ ;
7:  $\tau^* = (\text{reject}, \mathcal{A}) \leftarrow \mathcal{A}^{\mathcal{O}}(PK_F, \text{PP}, L_F, SK_{\mathcal{A}})$ ;
8:  $(\{EK_{F,S}\}, L_F) \leftarrow \text{Revoke}(\text{MK}, \tau^*, F, L_F)$ ;
9: for  $i = 1$  to  $n$  do
10:    $\{\sigma_{x_i^*}, VK_{F,x_i^*}, b_i\} \leftarrow \text{ProbGen}(\{t_i^*, x_i^*\}, PK_F)$ ;
11:    $\sigma_{y^*} \leftarrow \mathcal{A}^{\mathcal{O}}(PK_F, \text{PP}, L_F, \{\sigma_{x_i^*}, VK_{F,x_i^*}\}, \{EK_{F,S}\}, SK_{\mathcal{A}})$ ;
12:   if  $\exists i \in [n]$  s.t.  $((\tilde{y}, \tau_{\sigma_{y^*}}) \leftarrow \text{Verify}(\text{PP}, \sigma_{y^*}, VK_{F,x_i^*}, L_F))$  and  $((\tilde{y}, \tau_{\sigma_y}) \neq (\perp, (\text{reject}, \mathcal{A})))$  then
13:     return 1
14:   else
15:     return 0

```

Game 3 $\text{Exp}_{\mathcal{A}}^{\text{VindictiveS}} [\mathcal{PVC}_{\text{KDC}}, F, 1^\lambda]$:

```

1:  $\{t_i^*, x_i^*\}_{i \in [n]} \leftarrow \mathcal{A}(1^\lambda)$ ;
2:  $(\text{PP}, \text{MK}) \leftarrow \text{Setup}(1^\lambda)$ ;
3:  $(PK_F, L_F) \leftarrow \text{FnInit}(\text{PP}, \text{MK}, F)$ ;
4:  $SK_{\mathcal{A}} \leftarrow \text{Register}(\text{PP}, \text{MK}, \mathcal{A})$ ;
5:  $L_F \leftarrow \mathcal{A}(PK_F, \text{PP}, L_F, SK_{\mathcal{A}})$ ;
6:  $(EK_{F,\mathcal{A}}, L_F) \leftarrow \text{Certify}(\text{PP}, \text{MK}, F, L_F, \mathcal{A})$ ;
7: for  $i = 1$  to  $n$  do
8:    $\{\sigma_{x_i^*}, VK_{F,x_i^*}, b_i\} \leftarrow \text{ProbGen}(\{t_i^*, x_i^*\}, PK_F)$ ;
9:    $\tilde{S} \leftarrow \mathcal{A}^{\mathcal{O}}(PK_F, \text{PP}, L_F, \{\sigma_{x_i^*}, VK_{F,x_i^*}\}, EK_{F,\mathcal{A}}, SK_{\mathcal{A}})$  subject to Condition 1;
10:   $\sigma_{y^*} \leftarrow \mathcal{A}^{\mathcal{O}, \text{Compute}}(PK_F, \text{PP}, L_F, \{\sigma_{x_i^*}, VK_{F,x_i^*}\}, EK_{F,\mathcal{A}}, SK_{\mathcal{A}})$  subject to Condition 2;
11:  if  $\exists i \in [n]$  s.t.  $((\tilde{y}, \tau_{\sigma_{y^*}}) \leftarrow \text{Verify}(\text{PP}, \sigma_{y^*}, VK_{F,x_i^*}, L_F))$  and  $((\tilde{y}, \tau_{\sigma_y}) = (\perp, (\text{reject}, \tilde{S})))$  and  $(\perp \leftarrow \text{Revoke}(\text{MK}, \tau_{\sigma_y}, F, L_F))$  then
12:    return 1
13:  else
14:    return 0

```

thus reducing the size of the server pool and increasing the future reward for the malicious server. In Game 3 the adversary must (on lines 9 and 10) output an invalid result σ_{y^*} and the ID of a server \tilde{S} that it aims to cause to be revoked. It is provided with the standard oracle access on line 9 and on line 10 additionally with oracle access to **Compute** such that he can see outputs returned by honest servers (i.e. modelling the adversary submitting computation requests to the system himself), subject to the following constraints:

1. No query was made of the form $\mathcal{O}^{\text{Register}}(\text{PP}, \text{MK}, \tilde{S})$;
 2. As above but also no query was made of the form $\mathcal{O}^{\text{Compute}}(\sigma_{x_i^*}, EK_{F,\tilde{S}_i}, SK_{\tilde{S}})$;
- The adversary wins if the KDC believes that \tilde{S} returned \tilde{y} and revokes \tilde{S} .

Vindictive Manager and Blind Verification. In Appendix B we give two more security games for the manager model, namely the notions of Vindictive

Managers and Blind Verification. Vindictive Managers captures that a manager, being an intermediary in the verification process, may try to accept an incorrect answer in order to convince the client (or other recipient) of incorrect results – in a simple case, a vindictive manager could simply return the value of the “result” to the client. In the game, we provide the adversary with an encoded output, and require him to output an incorrect result, μ , with an acceptance token which will be accepted by a verifier in `VC.RetrieveOutput`. Thus, the goal is essentially to forge an encoded output μ .

The Blind Verification game captures a (weak) notion of output privacy in that it prevents verifiers from learning the output unless they hold an additional key, b . It does not prevent the servers themselves learning the output during `Compute` as usually considered output privacy. The challenger selects a random input x and gives the adversary the encoded output of the computation. The adversary must guess $F(x)$.

4 Construction

We provide a full construction of PVC-KDC using revocable KP-ABE in Appendix A. Informally, the scheme operates in the following way.

1. `VC.Setup` establishes public parameters and a master secret key by calling the `ABE.Setup` algorithm twice. In the manager model, we would require these algorithms to be run over the same set of random coins when choosing random exponents for attribute group elements such that the set of values for \mathcal{U} is the same in both cases, even if the associated semantic meaning differs. Thus, an adversary cannot recognize which set of parameters a given ciphertext belongs to, and hence cannot break the blind verification property. This algorithm also initializes a list of registered servers L_{Reg} and a time source τ^5 .
2. `VC.FnInit` initializes a list of servers L_F authorized to compute function F .
3. `VC.Register` creates a public-private key pair by calling the signature `KeyGen` algorithm. This is run by the KDC (or the manager in the manager model) and updates L_{Reg} to include S .
4. `VC.Certify` creates the key $EK_{F,S}$ that will be used by a server S to compute F by calling the `ABE.KeyGen` and `ABE.KeyUpdate` algorithms twice – once with a “policy” for F and once with the complement \bar{F} . The algorithm also updates L_F to include S .
5. `VC.ProbGen` creates a problem instance $\sigma_x = (c_0, c_1)$ by encrypting two randomly chosen messages, and a verification key $VK_{F,x}$ by applying a pre-image resistant hash function g to the messages. The ciphertexts and verification tokens are ordered randomly according to a bit b , such that the positioning of an element does not imply whether it relates to F or for \bar{F} .
6. `VC.Compute` is run by a server S and computes $F(x)$. Given a problem instance $\sigma_x = (c_0, c_1)$ it returns (m_0, \perp) if $F(x) = 1$ or (\perp, m_1) if $F(x) =$

⁵ τ could be a counter that is maintained in the public parameters or a networked clock.

- 0, ordered according to b chosen in VC.ProbGen , together with a digital signature computed over the output. The server can determine the value of b based on the results of decryptions with the different ABE parameters, and order his output correspondingly.
7. VC.Verify either accepts the output $\sigma_y = (d_0, d_1)$ or rejects it. This algorithm verifies the signature on the output and then confirms the output is correct by applying g and comparing with $VK_{F,x}$. In VC.BlindVerify the verifier can compare pairwise between the components of σ_y and $VK_{F,x}$ to determine correctness but as they are unaware of the value of b , they do not know the order of these elements and therefore do not learn whether the correct output corresponds to F or \bar{F} being satisfied i.e. if $F(x) = 1$ or 0 respectively. The verifier outputs an **accept** or **reject** token as well as the satisfying (non- \perp) output value $\mu \in \{d_b, d_{1-b}\}$. In VC.RetrieveOutput a verifier that has knowledge of b can check whether the output from BlindVerify matches m_0 or m_1 .
 8. VC.Revoke is run by the KDC and redistributes fresh keys to all non-revoked servers. This algorithm updates L_F and updates $EK_{F,S}$ using the results of two calls to the ABE.KeyUpdate algorithm.

Theorem 1. *Given a secure revocable KP-ABE scheme in the sense of indistinguishability against selective-target with semi-static query attack (IND-sHRSS) [1] for a class of functions \mathcal{F} closed under complement, a signature scheme secure against EUF-CMA and a pre-image resistant hash function g , let VC be the verifiable computation scheme defined in Algorithms 1–9. Then VC is secure in the sense of Public Verifiability, Revocation, Vindictive Servers, Blind Verification and Vindictive Managers.*

Informally, the proofs of Public Verifiability and against Vindictive Managers rely on the IND-CPA security of the underlying revocable KP-ABE scheme and the pre-image resistance of the function g . Revocation relies on the IND-sHRSS security of the revocable KP-ABE scheme. Blind Verification relies on the indistinguishability of two random messages, and the pre-image resistance of g . These proofs are left for the full version of the paper. Now we present a proof sketch for the security against vindictive servers.

Proof (Sketch). Let \mathcal{A}_{VC} be an adversary with non-negligible advantage against the Vindictive Servers game (Game 3). We construct an adversary \mathcal{A}_{Sig} with non-negligible advantage δ in the EUF-CMA signatures game using \mathcal{A}_{VC} . \mathcal{A}_{Sig} interacts with the challenger \mathcal{C} in the EUF-CMA security game and acts as the challenger for \mathcal{A}_{VC} in the security game for Vindictive Servers for a function F . Here the idea is that \mathcal{A}_{Sig} can create a VC instance and play the Vindictive Servers game with \mathcal{A}_{VC} by executing Algorithms 1–9 himself. \mathcal{A}_{Sig} will guess a server identity that he thinks the adversary will select to vindictively revoke. The signature signing key that would be generated during the **Register** algorithm for this server will be implicitly set to be the signing key in the EUF-CMA game and any **Compute** oracle queries for this identity will be forwarded to the challenger to compute. Then, assuming that \mathcal{A}_{Sig} guessed the correct server

identity, \mathcal{A}_{VC} will output a forged signature that \mathcal{A}_{Sig} may output as its guess in the EUF-CMA game. If \mathcal{A}_{Sig} guessed the challenge identity correctly (i.e. $\bar{S} = \tilde{S}$) then \mathcal{A}_{Sig} succeeds with the same non-negligible advantage δ as \mathcal{A}_{VC} . Let $n = |\mathcal{U}_{id}|$, then the probability that \mathcal{A}_{Sig} correctly guesses $\bar{S} = \tilde{S}$ is $\frac{1}{n}$ and $Adv_{\mathcal{A}_{Sig}} \geq \frac{1}{n} Adv_{\mathcal{A}_{VC}} \geq \frac{\delta}{n} \geq \text{negl}(\lambda)$. Thus we conclude that \mathcal{A}_{Sig} has a non-negligible advantage against the EUF-CMA game if \mathcal{A}_{VC} has a non-negligible advantage in the Vindictive Servers game, but as we assume the signature scheme in our construction to be EUF-CMA secure, such an adversary may not exist.

References

1. N. Attrapadung and H. Imai. Attribute-based encryption supporting direct/indirect revocation modes. In M. G. Parker, editor, *IMA Int. Conf.*, volume 5921 of *Lecture Notes in Computer Science*, pages 278–300. Springer, 2009.
2. H. Carter, C. Lever, and P. Traynor. Whitewash: Outsourcing garbled circuit generation for mobile devices. Cryptology ePrint Archive, Report 2014/224, 2014. <http://eprint.iacr.org/>.
3. S. G. Choi, J. Katz, R. Kumaresan, and C. Cid. Multi-client non-interactive verifiable computation. In *TCC*, pages 499–518, 2013.
4. R. Gennaro, C. Gentry, and B. Parno. Non-interactive verifiable computing: Outsourcing computation to untrusted workers. In T. Rabin, editor, *CRYPTO*, volume 6223 of *Lecture Notes in Computer Science*, pages 465–482. Springer, 2010.
5. C. Gentry. Fully homomorphic encryption using ideal lattices. In M. Mitzenmacher, editor, *STOC*, pages 169–178. ACM, 2009.
6. S. Goldwasser, V. Goyal, A. Jain, and A. Sahai. Multi-input functional encryption. Cryptology ePrint Archive, Report 2013/727, 2013. <http://eprint.iacr.org/>.
7. V. Goyal, O. Pandey, A. Sahai, and B. Waters. Attribute-based encryption for fine-grained access control of encrypted data. *IACR Cryptology ePrint Archive*, 2006:309, 2006.
8. B. Parno, M. Raykova, and V. Vaikuntanathan. How to delegate and verify in public: Verifiable computation from attribute-based encryption. Cryptology ePrint Archive, Report 2011/597, 2011.
9. A. C.-C. Yao. How to generate and exchange secrets (extended abstract). In *FOCS*, pages 162–167. IEEE Computer Society, 1986.

A Construction

We now provide an instantiation of a PVC-KDC scheme. Our construction is based on that used by Parno et al. [8] (summarised in Sec. 2) which uses Key-Policy Attribute-based Encryption (KP-ABE) in a black-box manner to outsource the computation of a Boolean function⁶. Notice that to achieve the outsourced evaluation of functions with n bit outputs, it is possible to evaluate

⁶ Following Parno et al. we restrict our attention to Boolean functions, and in particular the complexity class NC^1 which includes all circuits of depth $\mathcal{O}(\log n)$. This class includes common functions of interest such as AND, OR, NOT, equality and comparison operators, arithmetic operators and regular expressions.

n different functions, each of which applies a mask to output the single bit in position i .

Recall that if \perp is returned by the server then the verifier is unable to determine whether $F(x) = 0$ or whether the server misbehaved. To avoid this issue, we follow Parno *et al.* and restrict the family of functions \mathcal{F} we can evaluate to be the set of Boolean functions closed under complement. That is, if F belongs to \mathcal{F} then \bar{F} , where $\bar{F}(x) = F(x) \oplus 1$, also belongs to \mathcal{F} . Then, the client encrypts two random messages m_0 and m_1 . The server is required to return the decryption of those ciphertexts. Thus, a well-formed response satisfies the following:

$$(d_0, d_1) = \begin{cases} (m_0, \perp), & \text{if } F(x) = 1; \\ (\perp, m_1), & \text{if } F(x) = 0. \end{cases}$$

Hence, the client will be able to detect whether the server has misbehaved. We assume the existence of a *revocable KP-ABE scheme* for a class of functions \mathcal{F} that is closed under complement. Such a scheme defines the algorithms ABE.Setup, ABE.KeyGen, ABE.KeyUpdate, ABE.Encrypt and ABE.Decrypt. We also make use of a signature scheme with algorithms Sig.KeyGen, Sig.Sign and Sig.Verify and a pre-image resistant hash function g .

Then we construct a publicly verifiable computation scheme for the same class of functions comprising the algorithms VC.Setup, VC.FnInit, VC.Register, VC.Certify, VC.ProbGen, VC.Compute, VC.Verify and VC.Revoke). More formally, our scheme is defined by Algorithms 1–9.

Algorithm 1 VC.Setup

- 1: Let $\mathcal{U} = \mathcal{U}_{\text{attr}} \cup \mathcal{U}_{\text{ID}} \cup \mathcal{U}_{\text{time}}$
 - 2: $(MPK_{\text{ABE}}^0, MSK_{\text{ABE}}^0) \leftarrow \text{ABE.Setup}(1^\lambda, \mathcal{U})$
 - 3: $(MPK_{\text{ABE}}^1, MPK_{\text{ABE}}^1) \leftarrow \text{ABE.Setup}(1^\lambda, \mathcal{U})$
 - 4: $L_{\text{Reg}} = \epsilon$ (i.e. an empty list is created)
 - 5: Initialize τ
 - 6: $\text{PP} = (MPK_{\text{ABE}}^0, MPK_{\text{ABE}}^1, L_{\text{Reg}}, t)$
 - 7: $\text{MK} = (MSK_{\text{ABE}}^0, MSK_{\text{ABE}}^1)$
-

Algorithm 2 VC.FnInit

- 1: Set $PK_F = \text{PP}$
 - 2: Set $L_F = \epsilon$ (i.e. an empty list is created).
-

Algorithm 3 VC.Register

- 1: $(SK_{\text{Sig}}, VK_{\text{Sig}}) \leftarrow \text{Sig.KeyGen}(1^\lambda)$
 - 2: $SK_S = SK_{\text{Sig}}$
 - 3: $L_{\text{Reg}} = L_{\text{Reg}} \cup (S, VK_{\text{Sig}})$
-

Algorithm 4 VC.Certify

- 1: $t \leftarrow \tau$
 - 2: $SK_{\text{ABE}}^0 \leftarrow \text{ABE.KeyGen}(S, F, MSK_{\text{ABE}}^0, MPK_{\text{ABE}}^0)$
 - 3: $SK_{\text{ABE}}^1 \leftarrow \text{ABE.KeyGen}(S, \bar{F}, MSK_{\text{ABE}}^1, MPK_{\text{ABE}}^1)$
 - 4: $UK_{L_F, t}^0 \leftarrow \text{ABE.KeyUpdate}(L_F, t, MSK_{\text{ABE}}^0, MPK_{\text{ABE}}^0)$
 - 5: $UK_{L_F, t}^1 \leftarrow \text{ABE.KeyUpdate}(L_F, t, MSK_{\text{ABE}}^1, MPK_{\text{ABE}}^1)$
 - 6: Output: $EK_{F, S} = (SK_{\text{ABE}}^0, SK_{\text{ABE}}^1, UK_{L_F, t}^0, UK_{L_F, t}^1)$ and $L_F = L_F \cup S$
-

Algorithm 5 VC.ProbGen

```
1:  $t \leftarrow \tau$ 
2:  $(m_0, m_1) \xleftarrow{\$} \mathcal{M} \times \mathcal{M}$ 
3:  $b \xleftarrow{\$} \{0, 1\}$ 
4:  $c_b \leftarrow \text{ABE.Encrypt}(t, x, m_b, \text{MPK}_{\text{ABE}}^0)$ 
5:  $c_{1-b} \leftarrow \text{ABE.Encrypt}(t, x, m_{1-b}, \text{MPK}_{\text{ABE}}^1)$ 
6: Output:  $\sigma_x = (c_b, c_{1-b})$  and  $VK_{F,x} = (g(m_b), g(m_{1-b}), L_{\text{Reg}})$ 
```

Algorithm 6 VC.Compute

```
1: Input:  $EK_{F,S} = (SK_{\text{ABE}}^0, SK_{\text{ABE}}^1, UK_{L_F,t}^0, UK_{L_F,t}^1)$  and  $\sigma_x = (c_b, c_{1-b})$ 
2: Parse  $\sigma_x$  as  $(c, c')$ 
3:  $d_0 \leftarrow \text{ABE.Decrypt}(c, SK_{\text{ABE}}^0, \text{MPK}_{\text{ABE}}^0, UK_{L_F,t}^0)$ 
4:  $d_1 \leftarrow \text{ABE.Decrypt}(c', SK_{\text{ABE}}^1, \text{MPK}_{\text{ABE}}^1, UK_{L_F,t}^1)$ 
5: if  $d_0 \neq \perp$  or  $d_1 \neq \perp$  then
6:    $b=0$ 
7: else
8:    $b=1$ 
9:    $d_0 \leftarrow \text{ABE.Decrypt}(c, SK_{\text{ABE}}^1, \text{MPK}_{\text{ABE}}^1, UK_{L_F,t}^1)$ 
10:   $d_1 \leftarrow \text{ABE.Decrypt}(c', SK_{\text{ABE}}^0, \text{MPK}_{\text{ABE}}^0, UK_{L_F,t}^0)$ 
11:  $\gamma \leftarrow \text{Sig.Sign}((d_b, d_{1-b}, S), SK_S)$ 
12: Output:  $\sigma_y = (d_b, d_{1-b}, S, \gamma)$ 
```

Algorithm 7 VC.BlindVerify

```
1: Input:  $VK_{F,x} = (g(m_b), g(m_{1-b}), L_{\text{Reg}})$  and  $\sigma_y = (d_b, d_{1-b}, S, \gamma)$ 
2: if  $S \in L_F$  and  $(S, VK_{\text{Sig}}) \in L_{\text{Reg}}$  then
3:   if  $\text{Sig.Verify}((d_b, d_{1-b}, S), \gamma, VK_{\text{Sig}}) \rightarrow \text{accept}$  then
4:     if  $g(m_b) = g(d_b)$  then
5:       Output  $(\mu = d_b, \tau_{\sigma_y} = (\text{accept}, S))$ 
6:     else if  $g(m_{1-b}) = g(d_{1-b})$  then
7:       Output  $(\mu = d_{1-b}, \tau_{\sigma_y} = (\text{accept}, S))$ 
8:     else
9:       Output  $(\mu = \perp, \tau_{\sigma_y} = (\text{reject}, S))$ 
10: Output  $(\mu = \perp, \tau_{\sigma_y} = (\text{reject}, \perp))$ 
```

Algorithm 8 VC.RetrieveOutput

```
1: Input:  $VK_{F,x} = (g(m_b), g(m_{1-b}), L_{\text{Reg}})$ ,  $\sigma_y = (d_b, d_{1-b}, S, \gamma)$ ,  $b$ , and  $(\mu, \tau_{\sigma_y})$  where  $\mu \in \{d_b, d_{1-b}, \perp\}$ 
2: if  $\tau_{\sigma_y} = (\text{accept}, S)$  and  $g(\mu) = g(m_0)$  then
3:   Output  $\tilde{y} = 1$ 
4: else if  $\tau_{\sigma_y} = (\text{accept}, S)$  and  $g(\mu) = g(m_1)$  then
5:   Output  $\tilde{y} = 0$ 
6: else
7:   Output  $\tilde{y} = \perp$ 
```

Algorithm 9 VC.Revoke

```
1: if  $\tau_{\sigma_y} = (\text{reject}, S)$  then
2:    $L_F = L_F \setminus S$ 
3:   Refresha  $\tau$ 
4:    $t \leftarrow \tau$ 
5:    $UK_{L_F,t}^0 \leftarrow \text{ABE.KeyUpdate}(L_F, t, \text{MSK}_{\text{ABE}}^0, \text{MPK}_{\text{ABE}}^0)$ 
6:    $UK_{L_F,t}^1 \leftarrow \text{ABE.KeyUpdate}(L_F, t, \text{MSK}_{\text{ABE}}^1, \text{MPK}_{\text{ABE}}^1)$ 
7:   for all  $S' \in L_F, S' \neq S$  do
8:     Parse  $EK_{F,S'}$  as  $(SK_{\text{ABE}}^0, SK_{\text{ABE}}^1, UK_{L_F,t-1}^0, UK_{L_F,t-1}^1)$ 
9:     Update and send  $EK_{F,S'} = (SK_{\text{ABE}}^0, SK_{\text{ABE}}^1, UK_{L_F,t}^0, UK_{L_F,t}^1)$ .
10: else
11:   output  $\perp$ 
```

^a By refresh we mean, for example, increment τ if it is a counter

B Security Games

Game 4 $\text{Exp}_{\mathcal{A}}^{\text{VindictiveM}} [\mathcal{PVC}_{\text{KDC}}, F, 1^\lambda]:$	Game 5 $\text{Exp}_{\mathcal{A}}^{\text{BVerif}} [\mathcal{PVC}_{\text{KDC}}, F, 1^\lambda]:$
1: $(\text{PP}, \text{MK}) \leftarrow \text{Setup}(1^\lambda);$ 2: $(\text{PK}_F, L_F) \leftarrow \text{FnInit}(\text{PP}, \text{MK}, F);$ 3: $\text{SK}_S \leftarrow \text{Register}(\text{PP}, \text{MK}, S);$ 4: $(\text{EK}_{F,S}, L_F) \leftarrow \text{Certify}(\text{PP}, \text{MK}, F, L_F, S);$ 5: $(t^*, x^*) \leftarrow \mathcal{A}^{\mathcal{O}}(\text{PK}_F, L_F, \text{PP});$ 6: $(\sigma_{x^*}, \text{VK}_{F,x^*}, b) \leftarrow \text{ProbGen}((t^*, x^*), \text{PK}_F);$ 7: $\sigma_y \leftarrow \text{Compute}(\sigma_{x^*}, \text{EK}_{F,S}, \text{SK}_S);$ 8: $(\mu, \tau_{\sigma_y}) \leftarrow \mathcal{A}^{\mathcal{O}}(\text{PP}, \sigma_y, \text{VK}_{F,x^*}, \text{PK}_F, L_F);$ 9: $\tilde{y} \leftarrow \text{RetrieveOutput}(\mu, \tau_{\sigma_y}, \text{VK}_{F,x^*}, b);$ 10: if $(\tilde{y} \neq F(x^*))$ and $(\tilde{y} \neq \perp)$ then 11: return 1 12: else 13: return 0	1: $(\text{PP}, \text{MK}) \leftarrow \text{Setup}(1^\lambda);$ 2: $(\text{PK}_F, L_F) \leftarrow \text{FnInit}(\text{PP}, \text{MK}, F);$ 3: $\text{SK}_S \leftarrow \text{Register}(\text{PP}, \text{MK}, S);$ 4: $(\text{EK}_{F,S}, L_F) \leftarrow \text{Certify}(\text{PP}, \text{MK}, F, L_F, S);$ 5: $x \xleftarrow{\$} \text{Dom}(F);$ 6: $t \xleftarrow{\$} \tau;$ 7: $(\sigma_x, \text{VK}_{F,x}, b) \leftarrow \text{ProbGen}((t, x), \text{PK}_F);$ 8: $\sigma_y \leftarrow \text{Compute}(\sigma_x, \text{EK}_{F,S}, \text{SK}_S);$ 9: $\hat{b} \leftarrow \mathcal{A}^{\mathcal{O}}(\sigma_y, \text{VK}_{F,x}, \text{PP}, \text{PK}_F, L_F);$ 10: if $\hat{b} = F(x)$ then 11: return 1 12: else 13: return 0

Vindictive Manager. In Game 4 we capture security against vindictive managers attempting to provide the client with an incorrect answer, as discussed in Section 3.2. This is a natural extension of the Public Verifiability notion (Game 1) in the manager model. The adversary, on line 5, chooses a challenge input value x , and the server computes an encoded output of $F(x)$. The adversary is then provided the encoded output and verification key and must output an encoded output μ and an acceptance token. The challenger runs **RetrieveOutput** on μ to get an output value \tilde{y} , and the adversary wins if the challenger accepts this output and $\tilde{y} \neq F(x)$. We remark that manager model instantiations may vary depending on the level of trust given to the manager. A completely trusted manager may simply return the result to a client, whilst a completely untrusted manager may have to provide the full output from the server and the client performs the full **Verify** step as well (in this case, security against vindictive managers will reduce to Public Verifiability since the manager would need to forge a full encoded output that passes a full verification step). Here we consider a middle ground where the manager is semi-trusted but the clients would still like a final, efficient check.

Blind Verification. With Game 5, we aim to show that a verifier that does not know the value of b chosen in **ProbGen** cannot learn the value of $F(x)$ given the encoded output. The challenger chooses an input value, x , at random from the domain of F and a time period, and uses these to generate an encoded input. He runs **Compute** on this input and gives the encoded output and the verification key to the adversary who must output a guess for the value of $F(x)$. We require that \mathcal{A} does not make a query to the **Certify** oracle for the function F else he may use trial decryptions to compute b , and may not submit x to the **ProbGen** oracle. Similarly, \mathcal{A} may not use the **Compute** oracle for σ_x , but we assume he does not get this value and by the IND-CPA property of the ABE scheme we use, he may not guess a valid ciphertext for x .